

Appl. No. 09/610,580  
Amdt. Dated February 26, 2004  
Reply to Final Office Action of November 26, 2003

Attorney Docket No. 81866.0028  
Customer No. 26021

**Remarks**

This is in response to the final Office Action dated November 26, 2003, which was paper # 9 of the present application. Claims 1-20 are pending in the present application. No amendments to the claims are offered. Applicant offers amendments to the specification and to the drawings to address matters raised by the Examiner. Entry of these amendments is requested as complying with matters of form and not touching the substance of the application. Reconsideration and early favorable action are respectfully requested.

The Examiner has objected to the drawings. Further to the Examiner's requirements, applicant submits proposed changes to the drawings. The attached sheets of drawings include changes to FIGS. 1-3 indicated in red. These sheets replace the original sheets. Formal drawings are included within this mailing that incorporate these changes.

Applicant offers amendments to the specification to correspond to the proposed changes to the drawings and to otherwise address the Examiner's comments. In particular, applicant submits herewith a substitute specification, excluding the claims. The substitute specification does not include new matter. Both a marked up version, in which underlining indicates additions and strikethrough indicates deletions, and a clean version of the substitute specification are enclosed.

The Office Action rejects the pending claims as anticipated by or obvious over U.S. Patent No. 5,751,956 to Kirsch, either taken alone or in combination with the Mann patent cited in the first Office Action. Applicant submits that the Kirsch patent is not related to URL forwarding and neither teaches nor suggests the applicant's system.

The present application's system allows a user's web site to be moved to different physical servers at different IP addresses without having to change the

domain name system (DNS) records for each change affecting the user's URL and IP address. In the claimed URL forwarding process, an initial URL request is provided to the forwarding server at a first IP address specified in the domain name system. The forwarding server redirects the initial URL request to a second web server at a second IP address different from the first IP address. This is not true of the Kirsch patent's system. In the Kirsch patent's system, a user's request for a first URL associated with a first IP address is redirected away from that first IP address to the second IP address of the Kirsch patent's server so that the first URL request does not reach the IP address specified in the DNS. Only after the Kirsch patent's server counts the user's request does the Kirsch patent's server allow the first URL request to resolve to the address recorded in the DNS.

Thus the Kirsch patent does not teach the present invention because it does not describe URL forwarding. This is discussed in greater detail below. The pending claims further distinguish over the Kirsch patent, taken alone or in combination, because the Kirsch patent does not teach or allow user set up or updating of a URL forwarding function. The Kirsch patent does not teach a domain management interface that facilitates such set up or updating URL forwarding functions. This is also discussed in greater detail below.

In a typical URL forwarding transaction, a user enters a URL (i.e., a domain name) into a browser and requests a web page. The browser directly or indirectly (i.e., through a cache) initiates contact with a domain name system (DNS) server that translates the URL into a first IP address. Having resolved the URL into the first IP address, the user's request is directed to a first server at the first IP address. The first server analyzes the received URL and, if the URL is defined for forwarding, checks a file in an associated file server to determine the second IP address to which the user's request should be forwarded and redirects that request to the "forwarded" second IP address. In this way, the URL forwarding system

causes a URL to resolve to an IP address other than the IP address associated with the URL in the domain name system.

This aspect of URL forwarding is reflected in independent claim 1, which recites:

“a first web server adapted to receive a request for a first URL ... identifying a first IP address of the first web server according to a domain name system; ... a first file associated with the first URL accessed contains a second IP address of a first destination server, the first web server returning the second IP address as part of the message in response to the first URL request.”

By contrast, the Kirsch patent does not describe URL forwarding. The Kirsch patent describes a system for more efficiently tracking the sort of “click through” web site accesses associated with user interaction with advertising content on a served web page. The Kirsch patent’s system operates on a web page to be served from the URL of the Kirsch patent’s server. The web page at the URL for the Kirsch patent’s server has a hyperlink to advertising content on an advertising web server at a first URL. Normally, clicking on this advertising hyperlink would resolve the first URL to the first IP address for the advertising web server. The Kirsch patent’s system modifies the advertising hyperlink so that clicking on the hyperlink results in another access to the URL for the Kirsch patent’s web server. Kirsch patent, column 6, lines 47-56. The Kirsch patent’s server counts the access to the advertising hyperlink and then passes the original hyperlink’s first URL to the original advertising web server at that first IP address. Kirsch patent, column 6, lines 57-59.

Thus, in the Kirsch patent’s system, a user’s URL request for an advertising web server at a first IP address (defined by the DNS) is redirected instead to the

Kirsch patent's web server so that the access can be counted and then the URL request is directed back to the originally intended first IP address. Kirsch patent, column 6, lines 47-59. In essence, this "temporary capture" of the URL request and redirection of that request to a second IP address before causing the URL request to resolve to the first IP address is the opposite of what is defined in claim 1 of this application. Specifically, claim 1 recites:

"a first web server adapted to receive a request for a first URL ... identifying a first IP address of the first web server according to a domain name system; ... a first file associated with the first URL accessed contains a second IP address of a first destination server, the first web server returning the second IP address as part of the message in response to the first URL request."

In the Kirsch patent's system, the web server that receives the request for the first URL is not at the first IP address (defined in the DNS). Simply put, the Kirsch patent's system does not accomplish URL forwarding. Consequently, the Kirsch patent does not meet claim 1 and so does not teach the inventions of claim 1 or its dependent claims.

There are a number of other differences between the Kirsch patent and the present application's system. The application's system works with a user interface that allows a user to directly define a URL forwarding function or to change the characteristics of a URL forwarding function. This is facilitated by, for example, use of the domain management interface described at pages 9-13 of the application. This aspect of the application's system is reflected in claim 1 as follows:

"the second web server adapted to receive a request to alter the second IP address within the first file to modify the association between the first URL and the second IP address of the first destination server, wherein the request to alter the association between the first URL and the second IP address is

provided through a domain management interface having a process for authenticating a user's right to modify contents of the first file."

The Kirsch patent does not describe such a user-modifiable URL forwarding system or the domain management interface preferably used to facilitate those user modifications. Rather, the Kirsch patent's system automates its process of redirecting URL requests derived from web pages originally served by the Kirsch patent's web server. This is simple to accomplish since the purpose of the Kirsch patent's system is to ensure that the Kirsch patent's web server keeps an accurate count of the number of times disparate users access a certain type of web content. The web server 16, illustrated in FIG. 1, generates blocks 56 and 58, illustrated in FIG. 3. The redirection data is generated entirely within the server and is never altered by user input. Block 72 of FIG. 4 of the Kirsch patent generates the count that is the whole purpose of the Kirsch patent; it does not change a URL mapping.

Consequently, for this additional reason, the Kirsch patent does not teach or suggest the present invention. The other references of record do not address this deficiency of the Kirsch patent.

Still further, the Kirsch patent does not meet claim 1's requirement of "a domain management interface having a process for authenticating a user's right to modify contents of the first file." This aspect of claim 1 relates to the fact that a user can modify the URL forwarding instructions, for example by changing the IP address of the server to which the URL is redirected, through a browser or similar tool. Validating the user's rights to modify the contents of the first file (which contains the forwarding information) is needed to allow such a remote, user-directed modification. This is not done in the Kirsch patent's system. The only sort of validation that is performed in the Kirsch patent's system is to check whether a client machine has the right to access certain website content. See Kirsch patent,

Appl. No. 09/610,580  
Amdt. Dated February 26, 2004  
Reply to Final Office Action of November 26, 2003

Attorney Docket No. 81866.0028  
Customer No. 26021

column 8, lines 3-19. The "tampering" referred to in this passage of Kirsch relates to a client machine repeatedly accessing a website to generate "hits" or "counts" that might translate into undeserved advertising revenue. This aspect of Kirsch has nothing to do with verifying that a user has the rights to change the address to which a URL will be redirected.

For the additional reason that the Kirsch patent does not teach "a domain management interface having a process for authenticating a user's right to modify contents of the first file," claim 1 distinguishes over the Kirsch patent. None of the other references of record address this aspect of the claimed invention. As such, claim 1 and its dependent claims 2-20 distinguish over the Kirsch patent taken alone or in combination.

For all of the reasons discussed above, independent claim 1 and its dependent claims 2-20 distinguish over the art of record and are in condition for allowance.

In view of the foregoing, it is respectfully submitted that the application is in condition for allowance. Reexamination and reconsideration of the application, as amended, are requested.

If for any reason the Examiner finds the application other than in condition for allowance, the Examiner is requested to call the undersigned attorney at the Los Angeles, California telephone number (213) 337-6700 to discuss the steps necessary for placing the application in condition for allowance.

Appl. No. 09/610,580  
Amdt. Dated February 26, 2004  
Reply to Final Office Action of November 26, 2003

Attorney Docket No. 81866.0028  
Customer No. 26021

If there are any fees due in connection with the filing of this response, please charge the fees to our Deposit Account No. 50-1314.

Respectfully submitted,  
HOGAN & HARTSON L.L.P.

Date: February 26, 2004

By: 

William H. Wright  
Registration No. 36,312  
Attorney for Applicant

500 South Grand Avenue, Suite 1900  
Los Angeles, California 90071  
Phone: 213-337-6700  
Fax: 213-337-6701

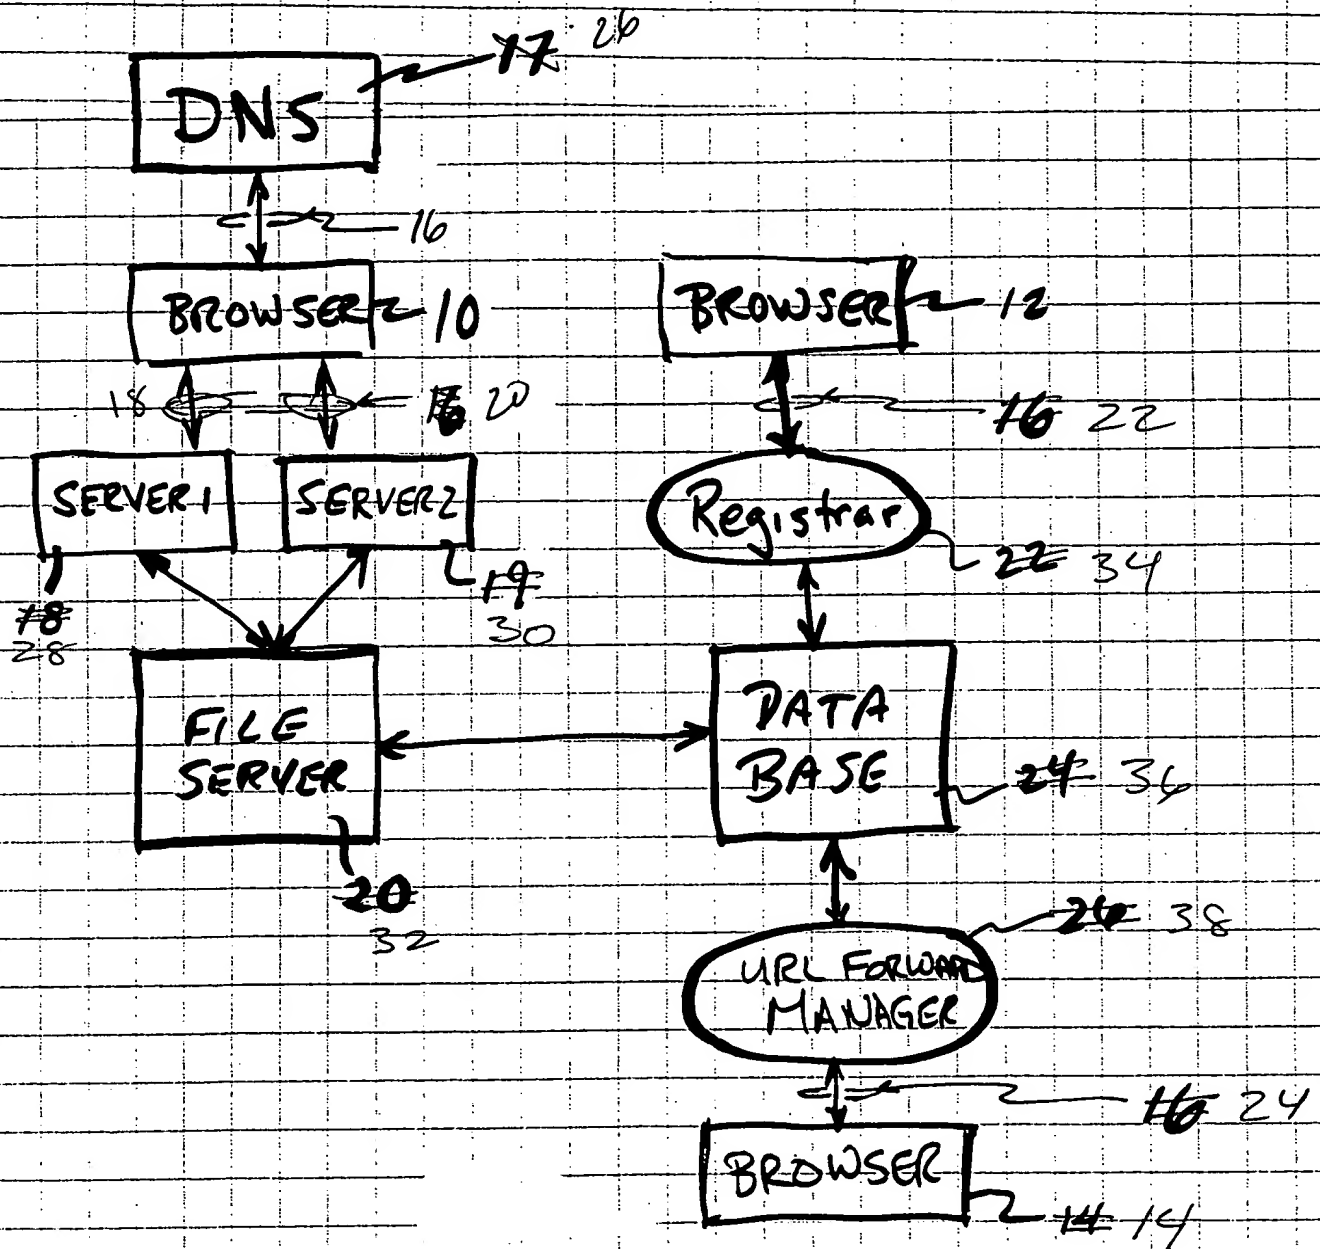




FIG. 1







Annotated Sheet Showing Changes




FIG. 3



09/610580

#11

SUBSTITUTE SPECIFICATION – VERSION MARKED TO SHOW CHANGES

PATENT  
Attorney Docket No. 81866.0028

5 **METHOD AND APPARATUS FOR URL FORWARDING**

**BACKGROUND OF THE INVENTION**

RECEIVED

MAR 04 2004

Technology Center 2100

1. Field of the Invention

10 [0001] The present invention relates generally to the Internet and more specifically to a method and apparatus for associating domain names or URLs with Internet protocol (IP) addresses and directing domain or URL requests to a desired IP address.

2. Description of the Related Art

15 [0002] Each computer on the Internet is identified by a unique Internet protocol (“IP”) address. This address is a 32-bit number organized as four 8-bit values separated by periods such as 123.45.67.89. Such a numerical system, while useful as a routing address system for computer-to-computer communication, is not human user-friendly. Consequently, domain names are used to allow users to more easily identify and connect to a target  
20 computer on the network. These user-friendly domain names (or “host names”), such as “register.com”, are easy for users to remember and, since they map to a unique IP number, accurately identify the computer’s IP address. In such a domain name identification scheme, the domain name forms a part of the uniform resource locator (URL) that specifies the location of resources on the World Wide Web. The URL identifies the mechanism used to  
25 access the resource (e.g., http, ftp, etc.), the specific computer that houses the resource, and the specific name of the resource (such as a filename).

[0003] As with the underlying Internet address, domain names typically have a hierarchical organization, with the trailing portion of the domain name, such as .com, .net, .org, .us, .uk or .jp, representing the top-level domain. Top-level domains include global top-  
30 level domains (gTLD) and country specific or country code top-level domains (ccTLD). The

global top-level domains include .com, .org, .net, .edu, .gov and .mil. Of these, .edu, .gov and .mil gTLD's are restricted to use by entities meeting specific qualifications. Country code top-level domains are country specific in that they identify registrations within a given country. The specific country governs registration for the country code top-level domains.

5 Some countries are "open" in that they allow any entity to register a domain name within its ccTLD. Other countries are "closed" and only allow entities that meet restrictions such as residency to register domain names in that ccTLD. Most domain users presently use one or more of the .com, .net or .org gTLDs.

[0004] The domain name entered by a user is sent over the Internet to a global  
10 network of servers called the "domain name system" (DNS), which receives the domain name as a request and translates the domain name into the target computer's numerical IP address. The numerical IP address is returned to the user's computer to enable it to connect to the target computer. Typically, after the user enters the domain name, the rest of the process is invisible to the user until the user connects to the target computer. The domain  
15 name system consists of a collection of root servers or DNS Servers that provide a directory linking domain names with corresponding IP addresses. There are presently thirteen root servers worldwide that contain authoritative databases listing all top-level domains. The collection of root servers is centrally managed for all global top-level domains to ensure that each computer on the network can be uniquely identified by unique domain names and  
20 numerical addresses.

[0005] A "registry" is an international organization or entity that is responsible for assigning domain names and Internet protocol addresses. Each country maintains its own registry, generally through a company or organization. The registry has the responsibility to record and update domain names and Internet protocol addresses, as well as the information  
25 associated with them, on the root servers. A registry is under contract from its respective government to control domain name registration. The registry may authorize other entities, known here as registrars, to conduct domain name registration and other aspects of the management of domain names and IP addresses.

[0006] A "registrar" is an organization or company that is authorized to provide  
30 registration services for all users of certain top-level domains, such as the .net, .org and .com global top-level domains. Registrars are presently authorized either by ICANN, the Internet

Corporation for Assigned Names & Numbers, a U.S. governmental organization under the Department of Commerce, or by the registrar's respective government to control domain name registration. A registrar is authorized by the registry to act as an agent of the registrar

to process domain name registration. The registrar has the responsibility to create and

maintain a Whois database and zone files for its customers. Examples of registrars presently include Register.com and Network Solutions, Inc., both authorized by ICANN.

[0007] A "registrant" is the individual or organization to whom a specific domain name is registered with the registry. Once a registrant has registered a domain name, paid the associated fees and met certain conditions, the individual or organization holds the domain name for use for a specific period of time. The registrant can use the domain name for such purposes as web hosting and e-mail. In many cases, the registrant may incorporate one or more domain names into an organizational identity or business. As such, a registration to use a particular domain name can be viewed as a significant asset for certain registrants.

[0008] The "shared registry system" (SRS) is a system that permits multiple registrars to provide registration services for the .com, .net and .org domains. The system is a shared database that holds information about domain names and their authoritative name servers. The shared registry system updates the root servers with information about the domain names within the .com, .org and .net gTLDs about every twenty-four hours in typical operation. The SRS allows accredited registrars to enter information about newly registered domain names into the SRS, and the information about the newly registered domain names is then uploaded to the root servers. Accredited registrars can update name server information within the SRS for domain names for which they are recognized as registrar. Accredited registrars are registered with the SRS and access the SRS through a secure and authenticated communication channel, such as through a secure socket level encrypted communication link. The SRS facilitates the updating of domain name and IP address information and also provides a utility for identifying the registrar that registered a domain name, when the entry to the SRS was created and the authoritative name servers for the domain name.

## SUMMARY OF THE PREFERRED EMBODIMENTS

[0009] An aspect of the present invention provides a URL processing system having a first web server, a file server and a second web server. The first web server is adapted to receive a request for a first URL and return a message associated with the first URL request. The file server is associated with and accessible by the first web server, the file server adapted to store a plurality of files corresponding to a plurality of URLs associated with a plurality of IP addresses. A first file is associated with the first URL and contains a first IP address of a first destination server, the first web server returning the first IP address as part of the message in response to the first URL request. The second web server is associated with the file server and is adapted to receive a request to alter the first IP address within the first file to alter the association between the first URL and the first IP address of the first destination server.

## BRIEF DESCRIPTION OF THE DRAWINGS

[0010] The embodiments and advantages of the present invention can be better understood in conjunction with the various drawings, which form a part of the disclosure of the present invention.

[0011] FIG. 1 illustrates schematically certain aspects of architecture for a preferred implementation of the present invention.

[0012] FIGS. 2(A) and 2(B) illustrates aspects of the process flow used in setting up URL forwarding in accordance with an implementation of aspects of the present invention.

[0013] FIG. 3 illustrates an exemplary process that might be used in an implementation of aspects of the present invention for modifying previously stored URL forwarding information.

## DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

[0014] Preferred embodiments of the present invention provide a method or apparatus for associating an IP address with a domain name and for readily altering the association

between the domain name and the IP address. In accordance with an embodiment of the present invention, a domain name is associated with a first web server within the domain name system (DNS) so that a request for a website or other resource associated with the domain name is presented to the first web server. The first web server accesses a file server to evaluate the requested domain name by accessing a file stored on the file server and determining from the contents of the file the IP address to which the domain name should resolve. The file associated with that domain name may contain information that initiates display of static or dynamic content from the first web server. Alternately, the file may contain instructions to associate the requested domain name with the IP address of a second web server that has the requested content or resource. Most preferably, the first web server performs this access to the file server directly, without executing a script or other interface program supplementary to the first web server. For example, the first web server may include a module that receives a domain name, accesses the file from the file server and evaluates the contents of the file. When the file includes an IP address to which the domain name should resolve, preferred implementations of the module recognize the presence of the IP address within the file and return the IP address of the second web server to the browser through which the user made the request. The browser can then use the returned IP address to access desired content from the destination web server.

**[0015]** Note here that the association between the domain name and the second, destination web server is maintained within a file server associated with the first, URL forwarding web server. This allows a domain name registrant or other URL user to readily alter the association between the domain name and an IP address without accessing the domain name system. In fact, two different associations exist: a first association between the domain name and a first IP address within the domain name system and a second association between the domain name and a second IP address. The first IP address is the address of the URL forwarding web server and the second IP address is the IP address of the second web server that stores the content associated with the domain name or URL. Preferably, the association between the URL and the IP address is maintained in a file within a file server that is directly accessed by the URL forwarding web server. By associating a domain name with the IP address of a destination server within a file stored on a file server accessed directly by the URL forwarding web server, the URL forwarding function is performed

quickly and efficiently. This architecture facilitates the easy initial association of a domain name with the IP address of a destination server that serves the content for that domain name and easy updating of the IP address associated with the domain name as the domain name registrant changes destination servers.

5   **[0016]**        Many companies seek to define their corporate identity at least in part in conjunction with one or more Internet domain names. This is true whether or not the company's business is conducted largely over the Internet. Building a brand identity in conjunction with a domain name makes a domain name registration valuable. When a company is successful in building the recognition of a domain name, it is important to  
10    maintain the domain name and to use that domain name as an important identifier for the company on the Internet.

**[0017]**        A company may initially provide its website on a low capacity server to take advantage of the lower cost and lower maintenance required by websites hosted in this manner. As the traffic to the company's website increases and the performance requirements  
15    for the website increase, higher capacity, higher performance servers are needed to accommodate the traffic and provide the desired performance. Evolving requirements cause companies to move through a succession of different website hosting options. For example, a company's initial website might consist of a modest home page hosted by an Internet service provider (ISP). A successor, higher traffic website might be hosted by a more  
20    substantial website hosting service. Eventually, the company might provide its own high traffic, high performance server or servers for its website. As a practical matter, a company may use many different website hosting options over the course of time in response to not only the performance concerns discussed above, but also to other concerns such as reliability and cost.

25   **[0018]**        Each different server is likely to be addressed through a different IP address. For the company's domain name to resolve to the correct IP address, the company's domain name needs to be properly associated with the server or servers that hosts the company's website. To do this, the association between the domain name and the host server's IP address is recorded in the domain name system (DNS) or in a system that is accessed through  
30    the DNS. The importance of maintaining a consistent domain name identity makes it important that this association be made quickly and reliably. Unfortunately, this is not



necessarily a task easily performed by those that are not expert in servers and the Internet or by those that do not have access to the resources and capabilities of a registrar. Aspects of the present invention are provided to facilitate one or more aspects of the process of properly associating a domain name or URL with a desired server at a defined IP address. In this way, a company can maintain a consistent identity and appearance on the Internet, despite changing the physical and Internet location of the servers that host the company's website or are associated with the company's domain name or names.

**[0019]** Preferred embodiments of the present invention provide implementations of a URL forwarding service that allows a URL to be mapped to an IP address simply and efficiently. This mapping is performed outside of the DNS so that the associations between the domain name and the IP address to which the domain name resolves can be altered easily. To be satisfactory, a URL forwarding service generally should be unobtrusive in that the forwarding or remapping operation should occur quickly, preferably with little noticeable delay in the user accessing the website being served the requested content. Consequently, it is desirable for a URL forwarding facility or method in accordance with the present invention to provide a prompt response to a URL request for content associated with a domain name.

**[0020]** One way in which a URL forwarding web server might be configured is through a relational database, such as an Oracle or similar database relating domain names to IP addresses. In such a configuration, a domain name is presented to a web server that accesses the relational database to retrieve the IP address corresponding to the domain name. To accomplish this, the web server receives the domain name and runs a common gateway interface (CGI) or script to interrogate the database, determines whether the domain name is to be redirected to a second, destination web server, and provides information to the web server in response to the request. The provided information might include the IP address of the second, destination server. While this configuration of URL forwarding web server is acceptable for a low volume application, this configuration has limited speed and does not scale well. As such, a more efficient and faster implementation of a URL forwarding server is preferred in accordance with the present invention.

**[0021]** FIG. 1 schematically illustrates an architecture that might be used in an embodiment of the present invention. Requests and other communications are made to the URL forwarding system from conventional web browsers 10, 12 and 14. FIG. 1 shows three

different browsers provided in the system, but these three different browsers could represent three different instances of the same browser. More practically, the browser 10 is likely to correspond to a visitor to a company's website, for example a customer of the company, and the browsers 12 and 14 are likely to represent a technical contact or other person associated with the company that maintains the company's websites and domain names. As the browsers 12 and 14 represent different ways of establishing and modifying the URL data stored within the file server, browsers 12 and 14 most likely represent different instances of a browser. In other regards, the illustrated browsers are conventional and each of the browsers 10, 12 and 14 preferably access the FIG. 1 system over ~~the Internet~~ connections 16, 18, 20, 22, 24. It should be recognized that much of the functionality accessed through the browsers 12 and 14 could be accessed through a more direct interface with the FIG. 1 system. Additionally, while the present discussion is provided with reference to the Internet, it should be appreciated that the various teachings herein are generally applicable to other communications networks that use similar addressing protocols.

[0022] The browser 10 issues a request, incorporating a URL, for a resource or content to the Internet 16, and this request is forwarded to one of the domain name servers of the domain name system (DNS) ~~17~~26. To illustrate aspects of the invention, this discussion assumes that the domain name or URL within the request issued by browser 10 is for a website or company that is using one of the services provided by the system of FIG. 1. For such domain names or URLs, the DNS ~~17~~26 returns to the browser 10 the IP address of the web servers ~~18, 19~~28, 30 shown in FIG. 1. The illustrated web servers are illustrated as comprising two different software web servers within two different hardware or physical servers. Those of ordinary skill in the art will recognize that such a configuration provides a higher traffic capacity and more uniform loading, as well as redundancy that facilitates higher reliability. In addition to the illustrated web servers, it is preferred that the servers ~~18, 19~~28, 30 include or operate in conjunction with a load balancing solution such as the Local Director products sold by Cisco Systems, Inc. (~~www.cisco.com~~). In accordance with a particularly preferred embodiment, the web servers ~~28, 30~~18, 19 are software web servers that can be customized by writing modules within the customizable web server's API. A preferred web server might, for example, be an HTTP/1.1 compliant web server such as that provided by the Apache HTTP Server Project of the Apache Software Foundation, which

presently can be found at and downloaded from www.apache.org that organization's web site. This open-sourced web server software is particularly preferred for its speed, readily available programming facilities and support, and its widespread acceptance and availability. The presently preferred Apache servers ~~18, 19~~ 28, 30 are modified in accordance with the present invention to include a module to handle the URL processing described here.

[0023] On receipt of the URL, the module within the web servers 28, 30 ~~18, 19~~ looks in the directory tree of the file server 2032 and retrieves the file associated with the URL request. For example, the URL request may comprise a domain name, the file name is most preferably made the same as the domain name and the directory structure may list the files, and hence the domain names, alphabetically. This allows the module to quickly identify and retrieve a file from the file server 2032 for a particular domain name. Most preferably, the file server provides high speed access to a large shared space of disk storage and may, for example, be a file server specifically adapted for serving data such as the file servers sold by Network Appliance, Inc. (www.netapp.com). Of course, with the evolution of technology, different file serving technology possibly not using disk drives will come into use. Aspects of the present invention are expected to benefit from such future technology.

[0024] If a file does not exist on the file server 2032 corresponding to the URL presented by the browser 10, then the module within the web servers 28, 30 ~~18, 19~~ returns a message to the browser 10 that the website is unavailable or that the website will be coming soon for that URL. Generally this message will be in the form of a single, static web page. Providing this sort of response allows a company to register its domain name and establish its URLs before actually completing construction of its websites. This gives an early stage company an opportunity to begin establishing an Internet presence and identity before it is even ready to launch a website.

[0025] It should be noted that files of zero length will be ignored within the file server and the file server 2032 will erase records or files within the file server 2032 by overwriting the files with length zero files. Consequently, the module is preferably programmed to ignore zero length files.

[0026] If a file exists on the file server 2032 corresponding to the URL, the module in the web servers will read the file to determine what services are associated with that URL. For example, the file may indicate that the URL is associated with a limited, generally static

web page stored on the file servers 28, 30 ~~18, 19~~ or on an associated web or file server. If that is the case, the module will cause that content to be returned to the browser 10.

[0027] The file might alternately indicate that URL forwarding is provided for that URL and return a message including at least an IP address through the web servers 28, 30 ~~18, 19~~ to the browser 10. Other information might also be provided to the browser. For example, if the URL forwarding web site wished to provide a free URL forwarding service by providing advertising, the web servers 28, 30 ~~18, 19~~ might return a framed web page, with the desired advertising information or other message provided on a peripheral part of the served, framed web page. Most preferably, the desired advertising or other information is provided in the frame such that the information has to be displayed on the user's display so long as the destination server content is displayed. Preferably the additional information is provided from a source other than the shared space on the file server 2032. When such a framed URL forwarding message is returned, the web servers 28, 30 ~~18, 19~~ cause the content from the destination server identified by the IP address retrieved from the file identified by the URL to be served onto the framed web page.

[0028] In other instances, the file on the file server 2032 identified by the URL received by the web servers 28, 30 ~~18, 19~~ might indicate that the URL's registrant has a higher level of URL forwarding service where the registrant or company pays for the service and the URL forwarding service is more transparent to the user of browser 10. In such a case the module within the web servers 28, 30 ~~18, 19~~ detects the type of service to be provided from the information in the file retrieved from the file server 2032 for that URL. The module provides the IP address information from the file server 2032 through the web server ~~18~~ 28 or ~~19~~ 30 and to the browser 10. This allows the domain name to be easily resolved into a desired IP address in a manner that allows a domain name to be associated with different servers or different IP addresses in a fairly simple manner.

[0029] This discussion now illustrates how a domain name registrant elects to use the URL forwarding service illustrated in FIG. 1 and how the files within file server 2032 are modified to reflect changes in the IP address to which the registrant's domain name should resolve. Preferred embodiments of the present invention might be used within a comprehensive domain name registration and domain management facility such as that described in U.S. patent applications Serial No. 09/560,433, filed April 27, 2000 and entitled

"Domain Manager and Method of Use," and Serial No. 09/587,403, filed June 5, 2000 and entitled Domain Manager for Plural Domains and Method of Use." These two pending U.S. patent applications are hereby incorporated by reference in their entirety and specifically for their teachings regarding interfaces through which a registrant enters information about a domain name including the information typically included in a Whois file and account information such as credit card numbers. The "registrar" web server 2234 indicated in FIG. 1 provides these types of facilities. In addition, these applications are incorporated by reference for their illustration of the information stored within a database such as the preferred Oracle database 2436 illustrated in FIG. 1.

**[0030]** The registrar web server 2234 is a facility used by individuals or organizations to register domain names and to manage various aspects of an Internet domain. Generally, the web server 2234 collects all of the information associated with a domain name and the domain name registrant. This information is all stored in the preferred Oracle database 2436. As discussed in the above-incorporated patent applications, registrar web server uses Perl or CGI programs or any of a number of different scripting languages to access and maintain the database 2436. Other information is preferably also collected and also stored within the database 2436. For example, the web server preferably collects information about different services the registrant wishes to use to facilitate domain management. Preferably, one of the options provided by the registrar web server 2234 is URL forwarding. For example, the registrant might be accessing the registrar web server 2234 from the browser 12. In such a case, a particular domain name will be active in the browser 12. The registrant, or the technical contact or other agent for the registrant, can simply select the URL forwarding functionality from a hyperlink displayed on the browser 12. The registrar web server 2234 functions to collect the additional information necessary to set up the URL forwarding functionality, including the initial IP address to which the domain name should resolve.

**[0031]** FIGS. 2(A) and 2(B) illustrates generally certain of the processes that might be employed in setting up URL forwarding, beginning from the main page 40 of the registrar web server. The process may include requesting 4042 a user to log in, preferably by accessing a login template 43. The process authenticates 44, 45 the user's rights to modify rights associated with the domain name. Although FIG. 2(A) shows a template 46 for indicating that the domain name is not registered through the registrar, this is not a

requirement for providing URL forwarding services. This is merely a matter of convenience and simplicity. The process may include various housekeeping steps such as getting pricing information for the URL forwarding functions from the database 2436 and getting the contact information associated with the domain name or other URL to be forwarded. After the main page 50 of the URL forwarding setup, the web server 2234 validates the credit card number provided by the registrant. The main page 50 may access an index template 54 to facilitate gathering or display of contact information. After the user's authority is validated 56, information is collected for the URL forwarding function. Instructions are assembled and confirmed 58, for example with reference to a confirmation template 60. The user's authority and financial information may be revalidated 62, 64. Eventually, the URL forwarding record 5466 is created, success is confirmed 68, a thank you template 72 is referenced to generate a thank you page 74 and the sequence of FIG. 2 is exited, 56 and the record 66 is written within the file stored to the file server 2032. In the event of errors in the processes of FIG. 2(B), the user is redirected to the main page 50, with reference to a reentry template 70 as appropriate.

[0032] When the information needed for the file is prepared, the database 2436 accesses a stored procedure and mounts the directory of the file server 2032. The stored procedure from database 2436 then writes to the file server 2032 an appropriate file including the desired information specifying the services to be provided and the IP address to which the URL should resolve.

[0033] Other database 2436 stored procedures can be used to perform maintenance on the file server 2032. For example, files for expired URL forwarding services can automatically be deleted by overwriting the files with length zero (blank) files after the expiration of their services. A different stored procedure might periodically delete empty files and defragment the disks of the file server 2032.

[0034] The FIG. 1 system preferably includes at least one mechanism for establishing a URL forwarding function and for altering the URL mapping once the URL forwarding is arranged. If URL forwarding is not set up through the registrar web server 2234, the FIG. 1 system preferably allows URL forwarding to be set up through a URL forwarding web server 2638. If no registrar web server 2234 is provided within the FIG. 1 system, then a URL forwarding web server 2638 is preferably provided that is capable of initially collecting the

information associated with a registrant and a domain name. In this regard, the URL forwarding web server 26-38 shares significant functionality with the registrar web server 2434.

[0035] The URL forwarding web server 26-38 is preferably used to modify the IP address to which a URL resolves or to modify any of the other information within the files stored on the file server 2032. A possible process flow for the URL modifying aspects of the server 26-38 is illustrated in FIG. 3, which is generally self-explanatory. Once a user is logged in 70-80, 84 and the user's rights are authenticated 7282, internal accounting information is passed 74-86 between the server 26-38 and the database 2436. The present status of the URL forwarding function, including the information within the corresponding file on the file server 2032, is obtained 76-88 and presented 78-90 on an interface to a user. This interface may query the user as to whether the user wants to modify or to delete the URL forwarding functionality. Depending on the option selected 80-92 by the user, the FIG. 3 process will proceed along the left hand path or the right hand path. The user chooses the left hand path to modify URL forwarding information by collecting a new forwarding URL, checking the validity of the URL 94, 96 and executing a change forwarding URL function 98. The user chooses the right hand path to delete a URL forwarding instruction, using a series of processes including confirming 100 the deletion instruction, deleting 102 the URL forwarding information from the file server 32 and displaying 104 the successful completion of the operation. The outcome of either path will cause the information in the file on the file server 20-32 to be altered following verification that the user wants the requested change implemented.

[0036] The illustrated process, and variations on the illustrated process, allow the IP address associated with a given URL to be modified in a prompt fashion. Because the DNS will continue to reference the URL to the IP address of the web servers 1828, 1930, this method effectively changes the URL to IP address mapping without accessing the DNS. Moreover, the preferred Apache web servers and file server architecture provide an efficient, high speed system that makes the use of URL forwarding both convenient and practical as generally being acceptable to customers from a performance point of view.

[0037] While aspects and certain advantages of the present invention have been described herein with reference to certain preferred embodiments of the present invention, it

should be appreciated that the present invention is not limited to the particular embodiments thereof. Those of ordinary skill in the art will appreciate that modifications and variations on the basic teachings of the present invention might be made without varying from the fundamental teachings thereof. Consequently, the scope of the present invention is to be  
5 determined from the claims, which follow.



## ABSTRACT OF THE INVENTION

A domain name is associated with the IP address of a first, URL forwarding web server within the domain name system (DNS) so that a request for a website or other resource associated with the domain name is presented to the URL forwarding web server. The URL forwarding web server accesses a file stored on a file server and determining from the contents of the file the IP address to which the domain name should resolve. The file associated with that domain name may contain information that initiates display of static or dynamic content from the URL forwarding web server. Alternately, the file may contain instructions to associate the requested domain name with the IP address of a second, destination web server that has the requested content or resource. Most preferably, the URL forwarding web server performs this access to the file server directly, without executing a script or other interface program supplementary to the URL forwarding web server. For example, the URL forwarding web server may include a module that receives a domain name, accesses the file from the file server and evaluates the contents of the file. When the file includes an IP address to which the domain name should resolve, preferred implementations of the module recognize the presence of the IP address within the file and return the IP address of the destination web server to the browser through which the user made the request.